

It Takes Advanced Microprocessor Technology to Make a Modern Datacenter

PART V OF A SERIES OF AMD WHITE PAPERS ON 64-BIT COMPUTING

Over fifty years ago, Herb Grosch, a scientist at IBM's Watson Research Center, observed that the cost of IBM's mainframe computers varied as the square root of their processing power.¹ Grosch noted that a system with four times the processing performance cost only twice as much, and modestly labeled his observation "Grosch's Law." The systems Grosch employed to calibrate his law used vacuum tubes, punched cards, and magnetic tape storage, since IBM did not adopt transistor technology until the late 1950s and IBM's San Jose Research Lab did not deliver its first magnetic disk storage system (the enormous 5MB RAMAC 305) until 1956. Nevertheless, IBM used Grosch's formulation as a key element in its mainframe pricing strategies through the 1960s and 1970s, thus driving customers seeking the most cost-effective computing solutions to acquire the largest systems they could afford. Since few applications needed the power of an entire mainframe, IBM developed a number of virtualization technologies that enabled users to combine separate applications and even disparate IBM operating system environments on a single machine. Sound familiar?

Forty years ago, Gordon Moore, then Director of Research and Development at Fairchild Semiconductor, first observed that the number of transistors on integrated circuits was doubling every 12 to 18 months, a phenomenon that has since become known as "Moore's Law." Early microprocessors were limited to a few thousand transistors and lacked serious computing capability. Evolving process technology and advanced processor architecture soon enabled designers to incorporate hundreds of millions (and even billions) of transistors into their chips so that within a few decades microprocessor performance challenged the capabilities of the largest mainframes. Moore's Law trumped Grosch's Law, and for many years smaller systems served as the least expensive means to deliver raw computing cycles. That shift drove the sales of almost one billion personal computers and tens of millions of x86-based servers over the past decade.

Today, many IT organizations struggle to remain afloat in a sea of servers, each of which is dedicated to running a single application. Large datacenters may contain tens of thousands of servers, each consuming a few hundred watts of power and emitting over a thousand BTUs per hour. Many servers in datacenters operate at less than fifteen percent of capacity and thus waste a large portion of the power they consume. Some analysts project that by 2010, 25 percent of all servers sold will go into datacenters with more than 100,000 servers.² Such centers will increasingly be located close to the least expensive sources of power, often near large hydroelectric generating facilities.

Just as the economics of Moore's Law shifted the center of gravity in the IT industry from larger, proprietary systems to smaller, industry-standard x86-based systems, the physics of powering and cooling systems—not to mention the logistics of installing, provisioning, and servicing them—in these massive server farms may very well nudge the pendulum back toward a smaller number of larger systems, or toward a combination that includes more virtualized four- and even eight-way x86-based servers and fewer standalone one- and two-way x86-based servers. In this document, we explore how a combination of dual- and quad-core processors, advanced virtualization technologies, and highly scalable four- and eight-way industry-standard servers may be used to create datacenters with fewer physical machines that use less power, generate less heat, and do more work, while simplifying the tasks of those who must build, install, operate, and maintain them. To the best of our knowledge, no one has quantified this phenomenon, although we know it when we see it. If someone finds a way to translate what is written here into the form of an equation, please speak up. You might see your name join Grosch, Moore, and Murphy in the canon of computing laws.

¹ InformationWeek, Apr 12, 2005 at <http://informationweek.com/story/showArticle.jhtml?articleID=160701576>

² Justin Ratner, IDF Keynote, September 26, 2006

HOW AMD TURNS LESS INTO MORE

AMD launched its first AMD Opteron™ processors in April 2003. Since that time, AMD has improved their performance by more than a factor of three, as they increased the operating frequency of the fastest cores from 1.4GHz to 3.0GHz, added a second processor core, and enhanced the system platforms that host these processors in ways both innovative and more obvious—like utilizing faster HyperTransport™ technology links and creating hardware-based AMD Virtualization™ technology (AMD-V™), or simply by moving to DDR2 DRAM technology, when that move made most sense for customers. The planned arrival of native Quad-Core AMD Opteron processors in mid-2007 will substantially increase the performance of AMD's processors yet again.

While AMD busied itself enhancing the AMD Opteron processor, the AMD Athlon™ 64 processor, and AMD Turion™ 64 mobile technology, its ecosystem partners focused on extending the range and flexibility of the system platforms that incorporate its processors. Customers can now purchase one- to four-way AMD Opteron processor-based platforms from HP, IBM, Sun, Dell, and many other system builder partners, while value-added resellers and integrators can tap into technology provided by infrastructure providers like Supermicro, Uniwide, and Tyan. Several vendors offer eight-way configurations as well. Equipped with AMD Opteron processors, these platforms allow customers to deploy affordable 16-core symmetric multiprocessing systems with performance that rivals that of the fastest mid-range proprietary and large SMP x86-based systems available just a few years ago. And soon, these eight-way systems will include 32 processing cores at prices similar to the 16-core systems you can buy today.

The underutilization of today's servers attests to the fact that most contemporary applications run well on two- and four-way servers; currently, only the most performance-hungry applications demand all the computational power that 16 or 32 AMD Opteron

processor cores can deliver. But just as virtualization technology enabled multiple applications to share a single IBM mainframe, the latest virtualization technologies from AMD software partners like Microsoft, Novell, Red Hat, Sun, SWSOFT, Virtual Iron, VMware, and XenSource can allow multiple virtual machines with multiple operating environments to share a single AMD Opteron processor-based server. A range of applications, hosted under a variety of environments such as the Windows®, Linux®, and Solaris operating systems, can coexist on a single system. Each application executes within its own container, or virtual machine (VM), unaware that it might in fact be one of many applications under one of many operating systems residing on a single system. Depending on the nature and duty cycle of the individually-hosted applications, one eight-way AMD Opteron processor-based system can handle as many as 80 VM environments.³ Since a typical eight-way AMD Opteron processor-based system might consume as much as 1.6KW at full load,⁴ each of these VMs could consume (on average) as little as 20 watts of power, a far cry from the still relatively efficient 200 watts that a real (as opposed to virtual) dual-processor, dual-core system might consume while running at a 10 percent utilization factor. These savings in real power consumption are amplified when you factor in the power needed (approximately 50 percent of required system power) to remove heat from the system: a 20 watt VM needs only a 10 watt HVAC system, while a 200 watt real machine can require up to 100 watts of power to run HVAC equipment.

Although the move from real to virtual environments can often be justified solely on the basis of reduced power consumption, the savings do not stop there. Without virtualization, the workloads running on 80 1U single- or dual-processor systems could occupy two 42U racks within the datacenter. Since a typical eight-way AMD Opteron processor-based system may require only 4U of rack space, the move from real machines to VMs could reduce the hardware footprint of system equipment within the datacenter by a factor

³ We assume here that a typical processor can support 10 virtual machines, when the workload running on a virtual machine might consume 10 percent of a real single-processor, dual-core server system.

⁴ Sun notes that an eight-CPU SunFire x4600 system consumes at most 1.603KW at full load.

of 20 or more. Even datacenters that have yet to max out power or HVAC resources can benefit from this dramatic reduction in the space required to handle aggregate workloads.

Those rare datacenters that still have power, HVAC, and space resources to spare also stand to benefit from a move to virtual environments. Virtualization software like VMware and the Xen open source hypervisor allow the rapid deployment of new VMs, along with the rapid relocation of VMs from one physical server to another. New applications can proceed from concept to development to deployment without the delays usually encountered when progress depends on an organization's ability to order, receive, and deploy incremental hardware resources. Many applications that run on real (i.e., non-virtual) machines require at least three dedicated systems in practice: one for development, one for normal execution, and one to back up the system normally used for execution. After an organization virtualizes its operations, it should need only two physical machines to accommodate its applications: one to support normal application execution and one to back up the first system and act as a development environment.

With fewer physical systems to manage, IT staff can be more productive and more responsive to user needs. And virtualization can also help enhance disaster recovery: should a VM fail for any reason, a new instance of a VM can be started at a remote location almost immediately.

SOUND TOO GOOD TO BE TRUE? WHAT'S THE CATCH?

Successful IT managers are usually able to distinguish real technological advances from the river of hyperbole that flows through their offices every day, and AMD realizes that no single solution is right for all customers. Indeed, its "customer-centric" philosophy mandates that it helps you identify a solution appropriate to meeting your needs.

First, AMD considers it important to note that the entire notion of virtual environments relies upon the assumption that any specific physical machine assigned to host VMs has available resources to support them—principally, compute cycles and physical memory. Systems using typical IT infrastructure applications usually meet this requirement because their utilization experiences peaks and valleys; systems deployed to support compute-hungry, high-performance technical computing applications may not. If a particular workload consumes most of a system's available resources on an on-going basis, it rarely pays to virtualize that workload.⁵

Second, AMD recognizes that the current state-of-the-art in virtualization technology imposes overheads and throughput constraints on most I/O operations. If a particular workload entails little CPU involvement but lots of I/O activity, virtualization may provide only limited benefit.⁶

And third, virtualization requires changes in the organizational mindset. In the "one-application, one-server" world, each department can clearly identify its computing assets. In a virtualized world, accounting may share a platform with both engineering and marketing. While this resource and cost sharing might demand a fundamental change in thinking, the reduced support costs, speedier deployment, and more responsive IT support should help overcome any mental hurdles.

What does all this mean? Moving to a world in which applications are freed from the bounds of physical hardware requires IT organizations to reassess the way they manage resources. The current one-server-per-application model is easy to understand and implement, although it often leads to the sub-optimal use of machine and power resources. Judiciously applied, the new model—one application per VM and many VMs per server—can improve resource utilization. For example, IT should work with users to

⁵ There still may be some benefit that accrues from rapid recovery from hardware failures, but there will be little if any benefit derived from improved asset utilization.

⁶ This may be a short term problem. AMD's roadmap includes new technologies that will improve virtual I/O performance in systems slated for availability in 2008.

identify which workloads can operate harmoniously on a single piece of hardware. When possible, it makes sense to combine on a single physical server a mix of VMs that lean most heavily on different mixes of system resources: memory, processor, and I/O bandwidth. System management tools need to monitor both physical machine and VM performance. They must also identify issues that can impact performance at both the physical and virtual levels. Additional tools might be required to check VMs for software licensing provisions and charge back purposes. Many of AMD's system and software ecosystem partners have already added features to their systems management roadmaps designed to facilitate the management of virtual environments. As IT practitioners shift their focus to managing virtual rather than physical resources, virtualization will emerge as a key architectural feature in the dynamic datacenter of the future.

FINALLY, A WORD FROM OUR SPONSOR.

If you've read this far, you'll have noticed there has been no mention of AMD Opteron™ processor-based servers out-virtualizing systems based on competitive processors. The reason for this is simple: detailed explanations of the exceptional virtualization features included in the Next-Generation AMD Opteron processors, known collectively as AMD Virtualization™ (AMD-V™), would detract from the greater points we hope to make.

In closing, here are a few key observations:

- The move from physical to virtual environments can have a dramatic impact on your datacenter's power, HVAC, and space requirements.
- Virtualization can facilitate your organization's ability to deploy new applications on a timely and cost-effective basis.
- The benefits of virtualization increase as you map more VMs onto a single physical server.
- Programs running in virtual environments tend to be less friendly to processor caching facilities than programs running in real (i.e., non-virtual) environments. Cache management strategies that work well in real environments may provide less benefit in virtual environments. (Corollary: The efficient access to main memory [DRAM] native to AMD64's Direct Connect Architecture is even more important in virtualized environments than in real ones.) ■



Advanced Micro Devices
One AMD Place
P.O. Box 3453
Sunnyvale, CA 94088-3453

www.amd.com